REPORT DOCUMENTATION PAGE

AFRL-SR-AR-TR-04-

0456

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing ins the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any oth reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis H Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	August 14, 2004	Final, 11/15/2000-5/14/2004

4. TITLE AND SUBTITLE	5. FUNDING NUMBERS
Integrated Simulation-Based Methodologies for Planning and Estimation	F496200110161

6. AUTHOR(S)

. Steven I. Marcus, Michael C. Fu, and Alan S. Willsky

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
Office of Research Administration and Advancement University of Maryland Lee Bldg Rm 3103 College Park, MD 20742	

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSORING / MONITORING AGENCY REPORT NUMBER
Air Force Office of Scientific Research 801 N. Randolph Street, Room 933 Arlington, VA 22203-1977	

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION / AVAILABILITY STATEMENT

Approved for public release; distribution unlimited

20040910 058

13. ABSTRACT *(Maximum 200 Words*

Significant progress was made in a number of proposed research areas. The first major task in the proposal involved incorporating simulation-based optimization (and, in particular, ordinal optimization) into dynamic optimization problems. In support of this task, progress was made on new sampling methods for Markov Decision Processes (MDPs), a new time aggregation approach for MDPs, simulation-based methods for weighted cost-to-go MDPs, approaches to proving the exponential convergence rate of ordinal comparisons, approximate receding horizon approaches to MDPs and Markov games, and new classes of stochastic approximation algorithms. In support of the second major task that involved estimation and control algorithms for dynamic hierarchical and graphical models, a variety of algorithms and analytical tools were developed for models on graphs with loops that exploit embedded loop-free structure. These algorithms offer the potential of significantly enhanced solutions to a variety of optimization problems critical to the Air Force. Another major task in the proposal involved risk-sensitive estimation and control. In support of this task, a new filtering scheme for the risk-sensitive state estimation of partially observed Markov chains was introduced and analyzed.

14. SUBJECT TERMS			15. NUMBER OF PAGES
Simulation, planning, estimation, optimization, Markov Decision Processes, graphical models			11
			16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102

FINAL REPORT

GRANT F49620-01-1-0161

NOVEMBER 15, 2000 – MAY 14, 2004

Principal Investigator: Steven I. Marcus
Co-Principal Investigators: Michael C. Fu, Alan S. Willsky

Office of Research Administration and Advancement
Rm. 2100 Lee Building
University of Maryland
College Park, MD 20742

1. Summary of Effort:

Significant progress was made in a number of proposed research areas. The first major task in the proposal involved incorporating simulation-based optimization (and, in particular, ordinal optimization) into dynamic optimization problems. In support of this task, we have made progress on new sampling methods for Markov Decision Processes (MDPs), a new time aggregation approach for MDPs, simulation-based methods for weighted cost-to-go MDPs, approaches to proving the exponential convergence rate of ordinal comparisons, approximate receding horizon approaches to MDPs and Markov games, and new classes of stochastic approximation algorithms.

In support of the second major task that involves estimation and control algorithms for dynamic hierarchical and graphical models, we have developed a variety of algorithms and analytical tools for models on graphs *with* loops that exploit embedded loop-free structure. These algorithms offer the potential of significantly enhanced solutions to a variety of optimization problems critical to the Air Force. Another major task in the proposal involved risk-sensitive estimation and control. In support of this task, we introduced and analyzed a new filtering scheme for the risk-sensitive state estimation of partially observed Markov chains.

2. Accomplishments/New Findings:

2.1 Incorporating Simulation Based Ordinal Optimization into Dynamic Optimization Problems

Based on recent results for multi-armed bandit problems, we proposed an "adaptive" sampling algorithm that approximates the optimal value of a finite horizon Markov decision process (MDP) with infinite state space but finite action space and bounded rewards. The algorithm adaptively chooses which action to sample as the sampling process proceeds, and it is proved that the estimate produced by the algorithm is asymptotically unbiased and the worst possible bias is bounded by a quantity that converges to zero at rate O ($H \ln N/N$), where H is the horizon length and N is the total number of samples that are used per state sampled in each stage. The worst-case running-time complexity of the algorithm is also analyzed. The algorithm can be used to create an approximate receding horizon control to solve infinite horizon MDPs.

We proposed a time aggregation approach for the solution of infinite horizon average cost Markov decision processes via policy iteration. In this approach, policy update is only carried out when the process visits a subset of the state space. As in state aggregation, this approach leads to a reduced state space, which may lead to a substantial reduction in computational and storage requirements, especially for problems with certain structural properties. However, in contrast to state aggregation, which generally results in an approximate model due to the loss of the Markov property, time aggregation suffers no loss of accuracy, because the Markov property is preserved. Single sample path-based estimation algorithms are developed that allow the time aggregation approach to be

implemented online for practical systems. Some numerical and simulation examples are presented to illustrate the ideas and potential computational savings.

We studied simulation-based algorithms for weighted cost-to-go Markov Decision Process (MDP) problems. We developed a two-timescale simulation-based gradient algorithm for weighted cost-to-go problems and prove its convergence. We illustrated the algorithm by carrying out numerical experiments, comparing it with two other algorithms in the literature. The numerical results indicate comparable convergence rates for a small example, so we discuss conditions under which the proposed two-timescale algorithm would be preferred due to implementation considerations.

Michael Fu, one of the leaders in the area of simulation for optimization, has updated his previous survey and tutorial papers on the subject and published a feature article on the subject.

The asymptotic exponential convergence rate of ordinal comparisons follows from well-known results in large deviations theory, where the critical condition is the existence of a finite moment generating function. We showed that this is both a necessary and sufficient condition, and also show how one can recover the exponential convergence rate in cases where the moment generating function is not finite. In particular, by working with appropriately truncated versions of the original random variables, the exponential convergence rate can be recovered.

We considered the solution of stochastic dynamic programs using sample path estimates. Applying the theory of large deviations, we established conditions under which the sample path optimal policy converges to the true optimal policy, for both finite and infinite horizon problems, at an asymptotically exponential convergence rate. This is in contrast with the usual canonical (inverse) square root rate associated with standard statistical output analysis for performance evaluation, here corresponding to estimation of the value (cost-to-go) function itself. These results have practical implications for Monte Carlo simulation-based solution approaches to stochastic dynamic programming problems where it is impractical to extract the explicit transition probabilities of the underlying system model. A portfolio selection problem in finance, interesting in its own right, is used to illustrate the convergence rate results.

Building on previous work on the receding horizon approach for solving Markov decision processes (MDPs), we analyze the performance of the *approximate* receding horizon approach in terms of infinite horizon average reward. In this approach, we choose a finite horizon and at each decision time, we solve the given MDP with the finite horizon for an approximately optimal current action and take the action to control the MDP. We then analyze a recently proposed on-line policy improvement scheme, called "rollout", by Bertsekas and Castanon, and a generalization of the rollout algorithm, "parallel rollout", in terms of the infinite horizon average reward in the framework of the (approximate) receding horizon control.

We have made progress in the analysis of various classes of stochastic approximation algorithms for simulation optimization. For example, we proposed and analyzed a new class of simultaneous perturbation stochastic approximation (SPSA) algorithms that are effective for high-dimensional simulation optimization problems. Extensive numerical experiments on a network of M/G/1 queues with feedback indicate that the deterministic sequence SPSA algorithms proposed in [x] perform significantly better than the corresponding randomized algorithms. In [xii], we consider Simultaneous Perturbation Stochastic Approximation (SPSA) for function minimization. The standard assumption for convergence is that the function be three times differentiable, although weaker assumptions have been used for special cases. However, all work that we are aware of at least requires differentiability. In this paper, we relax the differentiability requirement and prove convergence using convex analysis.

We have applied some of our methodologies to problems in finance to test their viablility. In particular, we applied some of our techniques for approximating stochastic control models to a problem in finance. In particular, we approximated the value function with a piecewise linear interpolation function, reducing the problem to a series of problems with known solutions. We provide two examples of finance problems where this approximation technique yields both upper and lower bounds on the true value.

2.2 Estimation and Control Algorithms for Dynamical Graphical Models

We have had several important developments in our work on estimation and optimization for hierarchical and graphical models. Motivated by earlier work on efficient algorithms for stochastic models on loop-free graphs (i.e., trees), we have developed a variety of algorithms and analytical tools for models on graphs *with* loops that exploit embedded loop-free structure. In particular, we have developed a new, powerful, and very promising family of what we are calling *tree-reparametrization (TRP) algorithms*. Each iteration of a TRP algorithm involves operations over a tree embedded in the graph. The critical idea here is the recognition that optimal estimation algorithms on trees correspond to performing a *refactorization* of the probability distribution for the entire process, one that explicitly exposes the marginal distributions for each node in the graph. TRP algorithms iterative perform this refactorization over a set of trees. In our work we have demonstrated that this algorithm has better convergence properties than previously developed algorithms and have also developed important theoretical results on the characterization of fixed points of these iterations, on necessary conditions for convergence using a pair of spanning trees, and on bounds on the errors in these algorithms. The latter results involve careful use of concepts in convex duality to obtain methods for optimizing our bounds over all embedded trees in a graph. Since there are generally many embedded trees, performing this optimization directly is completely intractable. However, the use of a dual formulation reduces this to a remarkably simple optimization problem. A paper on the basic TRP formulation received an honorable mention for best paper award at NIPS'01, while a paper on the optimized bounds just described received the best paper award at UAI'02.

The work just described on TRP focuses on computing the marginal probability distributions at each node in a graph, allowing one to compute optimal local estimates. However, there is another problem of great practical importance—e.g., in multisensor data association and in coding applications—namely that of computing the overall MAP estimate, i.e., the peak of the overall joint distribution for all variables on the entire graph. If the graph is a tree—i.e., contains no cycles—this computation can be performed very efficiently either in a two-sweep fashion, generalizing the dynamic programming structure of the celebrated Viterbi algorithm or in a local message-passing algorithm, often referred to as the *max-product algorithm*. However, if the graph of interest contains loops, performing MAP estimation is, in general NP-Hard, and the application of the max-product algorithm to such graphs not only may not converge but also leads to suboptimal solutions when convergence does occur. In our work we have developed counterparts of our TRP algorithms that are aimed at the MAP problem instead. In part this work allows us to develop both algorithms that converge more frequently than standard max-product algorithms and also analyses of the properties of fixed points. In addition, by adapting our ideas on using multiple trees (used in our TRP analysis as the basis for our optimized bounds) we have been able to develop a *tree-reweighting* algorithm that is guaranteed to converge to the true MAP estimate for a nontrivial and practically important set of problems. These algorithms offer the potential of significantly enhanced solutions to a variety of optimization problems critical to Air Force C2ISR applications, including the notoriously complex problem of multisensor, multitarget data association.

In work on multi-scale models that is closely related to our work on graphical models, we have proposed a simple analytical model of an M time-scale Markov Decision Process (MMDP) for hierarchically structured sequential decision making processes, where decisions in each level in the M-level hierarchy are made in M different time-scales. In this model, the state space and the control space of each level in the hierarchy are non-overlapping with those of the other levels, respectively, and the hierarchy is structured in a "pyramid" such that a decision made at level m (slower time-scale) state and/or the state will affect the evolutionary decision making process of the lower level $m + 1$ (faster time-scale) until a new decision is made at the higher level but the lower level decisions themselves do not affect the higher level's transition dynamics. The performance produced by the lower level's decisions will affect the higher level's decisions. A hierarchical objective function is defined such that the finite-horizon value of following a (nonstationary) policy at the level $m+1$ over a decision epoch of the level m plus an immediate reward at the level m is the single step reward for the level m decision making process. From this we define a "multi-level optimal value function" and derive a "multi-level optimality equation". We have studied how to solve MMDPs exactly or approximately and also have also studied heuristic on-line methods to solve MMDPs. Finally, we have studied a number of decision and control problems that can be modeled as MMDPs.

2.3 Robust and Risk Sensitive Estimation and Control

We have studied risk-sensitive estimation for the Hidden Markov Models from a dynamical systems point of view. We have shown that risk-sensitive estimators belong to a broader class of *product estimators* in which risk-sensitivity is related to certain *scaling functions*. The product structure and the scaling functions perspective result in new insights into the underlying mechanism of risk-sensitive estimation. For the first time, in a series of theorems and examples, we have relate risk-sensitivity to the dynamics of the underlying process and exposed relations among the transition probabilities, risk-sensitivity, and the decision regions.

We introduced a sequential filtering scheme for the risk-sensitive state estimation of partially observed Markov chains. Our risk-sensitive Maximum A Posteriori Probability (MAP) estimators generalize the previously studied risk-sensitive filters. Structural results, the influence of the availability of information, mixing and non-mixing dynamics and the connection with other risk-sensitive estimation methods are considered. A qualitative analysis of the sample paths clarifies the underlying mechanism.

3. Personnel Supported

      a. Faculty: Steven I. Marcus, Michael C. Fu, Alan S. Willsky
      b. Post-Docs: Andrew Lim, Grazyna Badowski, Hong-Gi Lee
      c. Graduate Students: Scott Laprise, Vahid Ramezani (graduate student through Dec. 2001, then post-doctoral researcher), Jiaqiao Hu, , Pedram Fard, Abraham Thomas, Martin J. Wainwright (graduate student through Jan. 2002, then post-doctoral researcher)

4. Publications

[i] H.S. Chang, M.C. Fu, and S.I. Marcus, "An Adaptive Sampling Algorithm for Solving Markov Decision Processes," to appear in *Operations Research*.

[ii] M. C. Fu, "Optimization for Simulation: Theory vs. Practice" (Feature Article), *INFORMS Journal on Computing*, forthcoming.

[iii] M.C. Fu and X. Jin, "Convergence of Simulation-Based Policies for Stochastic Dynamic Programming with Application to Finance," *Mathematics of Operations Research*, submitted, 2002.

[iv] H.S. Chang and S.I. Marcus, "On Approximate Receding Horizon Approach for Markov Decision Processes: Average Reward Case," *Journal of Mathematical Analysis and Applications*, 286, October 2003, 636-651.

[v] H.S. Chang and S.I. Marcus, "Two-Person Zero-Sum Markov Games: Receding Horizon Approach," *IEEE Transactions on Automatic Control*, 48, November 2003, 1951-1961.

[vi] S. Bhatnagar, M.C. Fu, S.I. Marcus, and I.-J. Wang, "Two-Timescale Simultaneous Perturbation Stochastic Approximation Using Deterministic Perturbation Sequences, " to appear in *ACM Transactions on Modeling and Computer Simulation*.

[vii] Y. He, M.C. Fu, and S.I. Marcus, "Two-Timescale Simulation-Based Policy Gradient Algorithm for Weighted Cost Markov Decision Processes," submitted to *Probability in the Engineering and Informational Sciences*.

[viii] Y. He, M. C. Fu, and S. I. Marcus, "Convergence of Simultaneous Perturbation Stochastic Approximation for Non-differentiable Optimization," submitted to *IEEE Transactions on Automatic Control*, 48, August 2003, 1459-1463.

[ix] X. Jin and M.C. Fu, "A Large Deviations Analysis of Quantile Estimation with Application to Value At Risk," *Management Science*, submitted, 2002.

[x] S.B. Laprise, M.C. Fu, S.I. Marcus, and A.E.B. Lim, "Pricing American-Style Derivatives with European Call Options," submitted to *Management Science*.

[xi] M.J. Wainwright, T. Jaakkola, and A.S. Willsky, "Tree-Based Reparametrization for Approximate Estimation of Stochastic Processes on Graphs with Cycles," *IEEE Trans. on Information Theory*, Vol. 49, No. 5, May 2003, pp. 1120-1146.

[xii] H.S. Chang, P. Fard, S.I. Marcus, and M. Shayman, "Multi-Scale Markov Decision Processes," submitted to *IEEE Transactions on Automatic Control*, 48, June 2003, 976-987.

[xiii] V. Ramezani and S.I. Marcus, "Estimation of Hidden Markov Models: Risk-Sensitive Filter Banks and Qualitative Analysis of Their Sample Paths," *IEEE Transactions on Automatic Control*, 47, December 2002, 1999-2009.

[xiv] V. Ramezani and S.I. Marcus, "Risk Sensitive Probability for Markov Chains," submitted to *Systems and Control Letters*.

[xv] H.S. Chang, H.G. Lee, M.C. Fu, and S.I. Marcus "Evolutionary Policy Iteration for Solving Markov Decision Processes" submitted to *IEEE Transactions on Automatic Control*.

[xvi] M.J. Wainwright, T.S. Jaakkola, and A.S. Willsky, "A New Class of Upper Bounds on the Log-Partition Function," submitted for publication to *IEEE Trans. on Information Theory*.

[xvii] M.J. Wainwright, T.S. Jaakkola, and A.S. Willsky, "Tree-Consistency and Bounds on the Max-Product Algorithm," in preparation.

[xviii] M.J. Wainwright, T.S. Jaakkola, and A.S. Willsky, "Exact MAP Estimates by Agreement on (Hyper)trees: Linear Programming and Message-Passing Approaches," submitted to *IEEE Trans. on Information Theory*.

[xix] L. Chen, M.J. Wainwright, M. Cetin, and A.S. Willsky, "Multitarge-Multisensor Data Association Using the Tree-Reweighted Max-Product Algorithm," presented at the Signal Processing, Sensor Fusion, and Target Recognition 12[th] Conf. at SPIE's 17[th] Aerosense Symposium, Orlando, FL, April 2003.

[xx] A.T. Ihler, E.B. Sudderth, W.T. Freeman, and A.S. Willsky, Efficient Multiscale Sampling from Products of Gaussian mixtures," submitted to NIPS'03.

[xxi] M. J. Wainwright, T.S. Jaakkola, and A.S. Willsky, "Tree-reweighted Belief Propagation Algorithms and Approximate ML: Estimation by Pseudomoment Matching," Workshop on Artificial Intelligence in Statistics, Jan. 2003.

[xxii] X.-R. Cao, Z. Ren, S. Bhatnagar, M. C. Fu, and S.I. Marcus, "A Time Aggregation Approach to Markov Decision Processes," *Automatica*, **38**, 2002, 929-943.

[xxiii] M. C. Fu, "Simulation Optimization," *Proceedings of the 2001 Winter Simulation Conference*, pp. 53-61.

[xxiv] M. C. Fu and X. Jin, "On the Convergence Rate of Ordinal Comparison of Random Variables," *IEEE Transactions on Automatic Control*, vol. 46, no. 12, Dec. 2001.

[xxv] S. B. Laprise, M. C. Fu, S. I. Marcus, and A. E. B. Lim, "A New Approach to Pricing American-Style Derivatives," *Proceedings of the 2001 Winter Simulation Conference*, pp. 329-337.

[xxvi] M.J. Wainwright, "Stochastic Processes on Graphs with Cycles: Geometric and Variational Approaches," Ph.D. Thesis, MIT, Jan. 2002.

[xxvii] M.J. Wainwright, T.S. Jaakkola, and A.S. Willsky, "Tree-Based Reparameterization for Approximate Estimation on Loopy Graphs," NIPS'01 (Runner-up—Student Best Paper Award).

[xxviii] M.J. Wainwright, T.S. Jaakkola, and A.S. Willsky, "A New Class of Upper Bounds on the Log-partition Function," Conf. on Uncertainty in Artificial Intelligence 2002 (Best Paper Award).

[xxix] V. Ramezani, "Product Estimators for Hidden Markov Models," Ph.D. Thesis, University of Maryland, Dec. 2001.

5. Interactions/Transitions:

a. Participation/presentations at meetings, conferences, seminars

M. C. Fu and X. Jin, "On the Convergence Rate of Ordinal Comparison of Random Variables," November 6, 2001, INFORMS National Meeting, Miami.

M. Fu, "Convergence Rate of Sample Path Optimal Policies for Stochastic Dynamic Programming with Application to Finance," July 25, 2001, INFORMS Applied Probability Conference, New York City.

M. Fu, "Pricing American-Style Options: A Comparison of Monte Carlo Approaches," September 11, 2001, Third IMACS Seminar on Monte Carlo Methods, Salzburg, Austria.

J. Chawla, M. Shayman, and S. I. Marcus, "Risk Sensitive Control of Semi-Markov Decision Processes," 2001 SIAM Control Conference, July 2001, San Diego.

S. Laprise, M. C. Fu, and S. I. Marcus, "A Quick Approximate Simulation Algorithm for Pricing American-Style Derivatives," 2001 Winter Simulation Conference, Washington, D.C., December 2001.

V. Ramezani and S. I. Marcus, "A Risk-sensitive Generalization of the Maximum A Posterior Probability Estimator for Hidden Markov Models," Stochastic Theory and Control, a workshop in honor of the 60th Birthday of Tyrone Duncan, October 18-20, 2001, Lawrence, KA.

M.J. Wainwright, T.S. Jaakkola, and A.S. Willsky, "Tree-Based Reparameterization for Approximate Estimation on Loopy Graphs," NIPS'01 (Runner-up—Student Best Paper Award).

M.J. Wainwright, T.S. Jaakkola, and A.S. Willsky, "A New Class of Upper Bounds on the Log-Partition Function," Snowbird Workshop on Learning, Snowbird, UT, April 2002.

M.J. Wainwright, T.S. Jaakkola, and A.S. Willsky, "Tree-Based Reparameterization for Approximate Estimation of Stochastic Processes on Graphs with Cycles," Int'l. Symp. on Information Theory, Lausanne, Switzerland, July 2002.

M.J. Wainwright, T.S. Jaakkola, and A.S. Willsky, "A New Class of Upper Bounds on the Log-partition Function," Conf. on Uncertainty in Aritificial Intelligence, August 2002 (Best Paper Award).

M. Fu, "Optimization via Simulation: Theory and Practice," University of Maryland, Scientific Computation Seminar, December 4, 2001.

J. Chen and M.C. Fu, "Efficient Sensitivity Analysis of Mortgage Backed Securities," 12th Annual Derivatives Securities Conference, New York City, April 27, 2002.

M. Fu, "An Adaptive Sampling Algorithm for Solving Markov Decision Processes," Cornell University, School of Operations Research and Industrial Engineering, April 30, 2002.

M. Fu, "A Large Deviations Analysis of Quantile Estimation with Application to Value at Risk," MIT Operations Research Center Seminar Series, May 9, 2002.

Y. He, M.C. Fu, and S.I. Marcus, "Convergence of Simultaneous Perturbation Stochastic Approximation for Nondifferentiable Optimization," Joint Statistical Meetings, American Statistical Association Meeting, New York City, August 14, 2002.

S. Laprise, M. C. Fu, and S. I. Marcus, "A Quick Approximate Simulation Algorithm for Pricing American-Style Derivatives," 2001 Winter Simulation Conference, Washington, D.C., December 2001.

b. Consultative and advisory functions

(1) Prof. Willsky has continued in his role as a member of the Air Force Scientific Advisory Board. In particular:
   a. He has twice served as a panel member for the AF/SAB S&T Review of AFRL/SN and the relevant parts of AFOSR supporting SN.
   b. He has twice served as a panel member for the AF/SAB S&T Review of AFRL/IF and the relevant parts of AFOSR supporting IF.
   c. He has participated in three AF/SAB summer studies. These each involved extensive visits to AF and other service organizations, and the writing of portions of the AF/SAB reports for the studies briefed to the Secretary of the Air Force and the AF Chief of Staff. In particular, during this past year, Prof. Willsky participated on the Information Integration and Management Panel for the AF/SAB study on Predictive Battlespace Awareness, commissioned directly by CSAF. Prof. Willsky's specific role in this study was to examine and make recommendations on S&T needs for fusion to support Predictive Battlespace Awareness (report is written and is currently under review by the AF/SAB prior to release)

(2) Prof. Willsky participated in the AFOSR-AFRL/IF Strategic Planning Workshop, held at Dartmouth College, August 2002. The intent of this meeting was to define strategic directions for research and development in the information sciences, broadly defined. Prof. Willsky will likely chair the next of these workshops, tentatively planned for 2004.

(3) Through contacts made on the AF/SAB, Prof. Willsky has been asked to act as an informal consultant to staff of the National Reconnaissance Office. In particular, Prof. Willsky has been asked to provide advice on future directions for information technology and fusion.

(4) Prof. Willsky has regularly acted as a consultant to Alphatech, Inc. in a number of research projects including ones that represent direct transitions of the technology being developed under our AFOSR Grant. He currently serves as Alphatech's Chief Scientific Consultant.

c. Transitions

(1) Transition of our graphical estimation methods to Alphatech for several programs for higher-level fusion (funded primarily by DARPA). The points of contact for this work are Dr. Mark Luettgen, and Dr. Eric Jones.

(2) Transition of our new algorithms for graphical optimization to several Alphatech programs, including new methods for data association for multitarget tracking (points of contact Dr. Robert Washburn and Dr. Mark Luettgen), to the extraction of "links" in complex, heterogeneous data and construction of models of behavior from huge data repositories under (points of contact Dr. Eric Jones and Dr. Robert Washburn), and to large-scale scheduling problems (point of contact Dr. Craig Lawrence).

6. New discoveries, inventions, or patent disclosures: None

7. Honors/Awards

- Michael Fu, Outstanding Systems Engineering Faculty Award, Institute for Systems Research, University of Maryland

- Martin Wainwright: Runner-up best student paper award, NIPS'01.

- Martin Wainwright, Tommi Jaakkola, Alan Willsky: Best paper award, UAI'02.

- Martin Wainwright: 2002 MIT EECS Sprowl Award for the Best thesis in computer science